UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/611,291	06/30/2003	R. Hugo Patterson	6368P002	2915

8791 7590 04/11/2008

BLAKELY SOKOLOFF TAYLOR & ZAFMAN
1279 OAKMEAD PARKWAY
SUNNYVALE, CA 94085-4040

EXAMINER	
STACE, BRENT S	
ART UNIT	**PAPER NUMBER**
2161	
MAIL DATE	**DELIVERY MODE**
04/11/2008	PAPER

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Office Action Summary	Application No.	Applicant(s)	
	10/611,291	PATTERSON, R. HUGO	
	Examiner	Art Unit	
	BRENT STACE	2161	

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE *3* MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

1)☒ Responsive to communication(s) filed on *31 January 2008*.

2a)☐ This action is **FINAL**. 2b)☒ This action is non-final.

3)☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

4)☒ Claim(s) *1,3-11,13-18,20-22,24-33,35-43 and 45-50* is/are pending in the application.

4a) Of the above claim(s) _____ is/are withdrawn from consideration.

5)☐ Claim(s) _____ is/are allowed.

6)☒ Claim(s) *1,3-11,13-18,20-22,24-33,35-43 and 45-50* is/are rejected.

7)☐ Claim(s) _____ is/are objected to.

8)☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

9)☐ The specification is objected to by the Examiner.

10)☒ The drawing(s) filed on *11 May 2006* is/are: a)☒ accepted or b)☐ objected to by the Examiner.

Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).

Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).

11)☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

12)☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).

a)☐ All b)☐ Some * c)☐ None of:

1.☐ Certified copies of the priority documents have been received.

2.☐ Certified copies of the priority documents have been received in Application No. _____.

3.☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

1) ☒ Notice of References Cited (PTO-892)
2) ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948)
3) ☐ Information Disclosure Statement(s) (PTO/SB/08) Paper No(s)/Mail Date _____.
4) ☐ Interview Summary (PTO-413) Paper No(s)/Mail Date. _____ .
5) ☐ Notice of Informal Patent Application
6) ☐ Other: _____.

DETAILED ACTION

Remarks

1. This communication is responsive to the amendment filed January 31st, 2008. Claims 1, 3-11, 13-18, 20-22, 24-33, 35-43, and 45-50 are pending. In the amendment filed January 31st, 2008, Claims 1, 8, 15, 20, 27, 31-33, 35-43, and 45-50 are amended, and Claims 1, 8, 15, 20, 27, 31, 33, 40, and 47 are independent. The examiner acknowledges that no new matter was introduced and the claims are supported by the specification.

Continued Examination Under 37 CFR 1.114

2. A request for continued examination under 37 CFR 1.114, including the fee set forth in 37 CFR 1.17(e), was filed in this application after final rejection. Since this application is eligible for continued examination under 37 CFR 1.114, and the fee set forth in 37 CFR 1.17(e) has been timely paid, the finality of the previous Office action has been withdrawn pursuant to 37 CFR 1.114. Applicant's submission filed on 1/31/2008 has been entered.

Response to Arguments

3. The Applicant's arguments filed January 31st, 2008 with respect to Claims 1, 3-11, 13-18, 20-22, 24-33, 35-43, and 45-50 have been considered but are either not persuasive or are moot in view of the new ground(s) of rejection. See below for details.

4. The examiner would like to note that Zwilling shrinking files before applying Hitz to take (any iteration of) a snapshot of the shrunken files is the correct paradigm for considering the Zwilling and Hitz references in combination.

5. As to the applicant's arguments with respect to Claims 1, 8, 15, 27, 31, 33, and 40 for the prior art(s) allegedly not teaching "each storage tree representing active snapshots of user-generated data stored from each one of a plurality of remote computer systems taken at different points in time," the examiner respectfully disagrees

Hitz teaches "each storage tree representing active snapshots of user-generated data … taken at different points in time" since Hitz creates snapshots of a file system (storage tree) (each storage tree representing active snapshots). Hitz's data is generally concerned with user-generated data such as documents in a file system that were told (at some point) to be stored by a user (Hitz, col. 4, lines 1-8). Also, hitz, teaches that the snapshot "is a read-only copy of an entire file system at a given instant when the snapshot is created" (Hitz, col. 17, lines 41-43). Since, Hitz, initially supports up to 20 snapshots, (Hitz, col. 4, lines 32-36) and each is for an instant when the snapshot is created, it appears that Hitz teaches the "taken at different points in time" limitation.

Zwilling teaches "stored from each one of a plurality of remote computer systems." The Zwilling reference is clear at Zwilling, col. 3, lines 15-23 that the Zwilling invention "may be practiced with other computer configurations, including...network PC's…The invention may also be practiced in distributed computing environments where tasks are performed by remote processing devices that are linked through a

communications network…program modules may be located in both local and remote memory storage deivces."

6. As to the applicant's arguments with respect to Claims 1, 27, 31, and 40 for the prior art(s) allegedly not teaching "the plurality of nodes including leaf nodes representing data blocks of the user-generated data stored from each one of the plurality of remote computer systems and root and intermediate nodes representing data blocks of system-generated data created by the storage device, the system-generated data including pointers to reference other data blocks of the user-generated data to enable overlapping and sharing of the user-generated data between and among different snapshots of the plurality of snapshots," the examiner respectfully disagrees. It should be noted that user-generated data and system-generated data can be construed to mean the same type of data. This is because, for instance, the user can have a text document that they made (user-generated). When they go about saving the document, they tell the storage system to save the document. The storage system then goes about saving the user-generated document. Since the storage system stores it though, and not the user, it can be seen as being system-generated data. Hitz teaches about typical file systems today where documents are stored within directories (or folders) in a hierarchical tree format (Hitz, col. 8, lines 40-60). This establishes leaf nodes of user-generated data and a root and intermediate nodes "representing data blocks of system-generated data created by the storage device." Additionally, Hitz includes pointers through the snapshot system to reference the data that was snapshotted "because a newly created shotshot tree references exactly the same blocks as the root inode…over

time, the snapshot and the active file system share fewer and fewer blocks." (Hitz, cols. 17-18, lines 66-16). As can be seen from this teaching, snapshot "share" blocks. As such, it appears that Hitz fully enables "overlapping and sharing of the user-generated data between and among different snapshots of the plurality of snapshots."

Again, Zwilling, as discussed above, teaches the "stored from each one of the plurality of remote computer systems" limitation.

7. As to the applicant's arguments with respect to Claims 1, 27, 31, and 40 for the prior art(s) allegedly not teaching "storage trees containing a plurality of nodes including lea[f] nodes of data generated by users on a plurality of remote computer systems," the examiner respectfully disagrees. This argument appears to have been met above.

8. As to the applicant's arguments with respect to Claims 8, 20, 33, and 47 for the prior art(s) allegedly not teaching "wherein each block of data is stored within a log as contiguous data in a sequential order and wraps around to the beginning of the log once the end is reached so that the size of the log does not increase beyond its originally created size," the examiner respectfully submits that this argument is moot in view of the new ground(s) of rejection below.

9. The other claims argued merely because of a dependency on a previously argued claim(s) in the arguments presented to the examiner, filed January 31st, 2008, are moot in view of the examiner's interpretation of the claims and art and are still considered rejected based on their respective rejections from at least a previous Office action (part(s) of recited again below).

Response to Amendment

Specification

10. The lengthy specification has not been checked to the extent necessary to determine the presence of all possible minor errors. Applicant's cooperation is requested in correcting any errors of which applicant may become aware in the specification.

Claim Objections

11. Claims 20, 27, 31, and 40 are objected to because of the following informalities:

a. Claim 20, line 7 ends with improper punctuation. The line recites “created size,; and”. This appears to be a grammatical error.

b. Claim 27 recites "system generated data" on line 8. there appears to be a missing "-" between “system” and “generated."

c. Claim 27 recites the limitation "the user-generated target data" in lines 9-10. There is insufficient antecedent basis for this limitation in the claim.

d. Claim 31 recites the limitation "the user-generated target data" in line 13. There is insufficient antecedent basis for this limitation in the claim.

e. Claim 40 recites the limitation "the user-generated target data" in lines 14-15. There is insufficient antecedent basis for this limitation in the claim.

f. Claim 40 recites “descendant” in line 24 while other parts of the claim recite “descendent.” This appears to be a typographical error.

Appropriate correction is required.

Claim Rejections - 35 USC § 103

12. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

> (a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negatived by the manner in which the invention was made.

13. Claims 1, 3-7, 15-18, 27-33, 35-39, 40-43, 45, and 46 are rejected under 35 U.S.C. 103(a) as being unpatentable over U.S. Patent No. 6,249,792 (Zwilling et al.) in view of U.S. Patent No. 5,963,962 (Hitz et al.)

For **Claim 1**, Zwilling teaches: "A method of garbage collecting in a storage device [Zwilling, col. 5, lines 29-52 with Zwilling, Fig. 2] comprising:

- locating blocks of data in a log that are both referenced by at least one other block of data stored … and within a range at a tail of the log using pruned walking, the range representing an address range within an allocated segment of the log, [Zwilling, col. 12, lines 24-53 with Zwilling, Fig. 2 with Zwilling, col. 8, lines 49-51 with Zwilling, col. 8, lines 60-64 with Zwilling, Figs. 3C and 3E]
- …stored from each one of a plurality of remote computer systems [Zwilling, col. 3, lines 15-23]…the plurality of nodes including leaf nodes representing data blocks of the user-generated data [Zwilling, col. 8, lines 49-64] stored from each one of the plurality of remote computer systems [Zwilling, col. 3, lines 15-23] and

root and intermediate nodes representing data blocks of system-generated data created by the storage device, [Zwilling, col. 8, lines 49-64]

- …copying the blocks of data that are referenced…and within the range to an unallocated segment of the log [Zwilling, col. 5, lines 34-52 with Zwilling, col. 12, lines 41-53]
- …marking the range at the tail of the log as unallocated so that at least a portion of an address space within the range can be reclaimed" [Zwilling, col. 5, lines 34-52 with Zwilling, col. 12, lines 31-34].

Zwilling discloses the above limitations but does not expressly teach:

- "…within a set of one or more storage trees
- …each storage tree representing active snapshots of user-generated data…taken at different points in time and having a plurality of nodes, each node representing a block of data of an active snapshot associated with each storage tree, … the system-generated data including pointers to reference other data blocks of the user-generated data to enable overlapping and sharing of the user-generated data between and among different snapshots of the plurality of snapshots;
- …by one or more other blocks of data of other nodes…wherein blocks of data that are not referenced by other blocks of data and within the range remain untouched."

With respect to Claim 1, an analogous art, Hitz, teaches:

- "…within a set of one or more storage trees [Hitz, cols. 17-18, lines 66-16]

- …each storage tree representing active snapshots of user-generated data [Hitz, cols. 17-18, lines 66-16] … taken at different points in time [Hitz, col. 17, lines 40-43] and having a plurality of nodes, [Hitz, cols. 17-18, lines 66-16 with Zwilling, col. 8, lines 49-64] each node representing a block of data of an active snapshot associated with each storage tree, [Hitz, col. 17, lines 40-49 with Hitz, cols. 17-18, lines 66-16 with Zwilling, col. 8, lines 49-64] …the system-generated data including pointers to reference other data blocks of the user-generated data to enable overlapping and sharing of the user-generated data between and among different snapshots of the plurality of snapshots [Hitz, cols. 17-18, lines 66-14 with Hitz, col. 18, lines 35-38];
- …by one or more other blocks of data of other nodes [Hitz, col. 18, lines 35-38] …wherein blocks of data that are not referenced by other blocks of data and within the range remain untouched" [Hitz, col. 15, lines 55-57 with Hitz, col. 16, lines 15-17 with Hitz, col. 20, lines 25-35 with Zwilling, col. 5, lines 29-52 with Zwilling, col. 6, lines 42-49 with Zwilling, col. 12, lines 36-37].

It would have been obvious to one of ordinary skill in the art at the time of invention having the teachings of Hitz and Zwilling before him/her to combine Hitz with Zwilling because both inventions are directed towards conserving file space (Zwilling shrinks and Hitz snapshots data that was changed at a block level (not a fie level)).

Hitz's invention would have been expected to successfully work well with Zwilling's invention because both inventions use file systems on computers. Zwilling discloses an on-line dynamic file shrink facility (title) comprising shrinking log files.

However, Zwilling does not expressly disclose storage trees each representing a snapshot and where unreferenced data remains untouched. Hitz discloses a write anywhere file-system layout (title) comprising storage trees (file system snapshot trees) each representing a snapshot and where unreferenced data remains untouched.

It would have been obvious to one of ordinary skill in the art at the time of invention having the teachings of Hitz and Zwilling before him/her to take the write anywhere file-system layout and the storage trees of snapshots from Hitz and install it into the invention of Zwilling, thereby offering the obvious advantage of taking snapshots of garbage collected/shrunken data, thus saving space. Shrinking the files prior to snapshoting them frees more space in Hitz thus creating more room in Hitz so that another later snapshot won't be prematurely forced out when space in Hitz runs out (by snapshots consuming unacceptable numbers of disk blocks). Also, shinking prior to snapshoting guarantees that Hitz will not back up empty blocks (since they are dirtied by becoming empty in a snapshot). This also conserves space in Hitz. In fact, Hitz desires files to be in a shrunken form (Hitz, col. 5, lines 37-40). Therefore, Hitz is merely using Zwilling's shrinking technique to obtain files that have no fragments.

Different embodiments of Zwilling are use in the rejection for Claim 1 and its respective dependant claims. Zwilling (as modified by Hitz) teaches all of Claim 1 with assistance from Zwilling's different embodiments.

Claim 3 can be mapped to Zwilling (as modified by Hitz) as follows: "The method of claim 1, wherein locating the blocks of data that are referenced and within the range includes determining a minimum value among addresses of descendent nodes of a

node, [Zwilling, cols. 8-9, lines 46-14] wherein the minimum value represents a minimum address offset of a node that is closest referenced from the blocks of data" [Zwilling, cols. 8-9, lines 46-14 with Zwilling, col. 5, lines 29-52 with Zwilling Fig. 2].

Claim 4 can be mapped to Zwilling (as modified by Hitz) as follows: "The method of claim 3, wherein a location table includes an entry for nodes that reference other nodes [Zwilling, col. 8, lines 46-64 with Zwilling, col. 10, lines 5-13] and wherein determining the minimum value among addresses of descendent nodes of the node includes retrieving the minimum value from an entry in the location table associated with the node" [Zwilling, cols. 8-9, lines 46-14 with Zwilling, col. 10, lines 5-13].

Claim 5 can be mapped to Zwilling (as modified by Hitz) as follows: "The method of claim 4, wherein locating the blocks of data that are referenced and within the range includes processing the descendent nodes of the node upon determining that the minimum value among the addresses of the descent nodes is within the range" [Zwilling, col. 5, lines 29-52].

Claim 6 can be mapped to Zwilling (as modified by Hitz) as follows: "The method of claim 5 comprising modifying the addresses of the copied blocks of data that are stored in the location table based on the new locations of the copied blocks of data in the log" [Zwilling, Fig. 3C with Zwilling, col. 8, lines 29-33 with Zwilling, col. 9, lines 4-7].

Claim 7 can be mapped to Zwilling (as modified by Hitz) as follows: "The method of claim 5 further comprising modifying the minimum value in the entry in the table associated with the node when the minimum value changes based on the new locations of the copied blocks of data that are associated with descendent nodes of the node"

[Zwilling, Fig. 3C with Zwilling, col. 8, lines 29-33 with Zwilling, col. 9, lines 4-7 with Zwilling, col. 5, lines 29-52 with Zwilling, col. 10, lines 5-13].

For **Claim 15**, Zwilling teaches: "A method of garbage collecting in a storage system [Zwilling, col. 5, lines 29-52 with Zwilling, Fig. 2] comprising."

With respect to Claim 15, Zwilling teaches in a different embodiment:

- "…performing following operations until each block of data that is active in a range to be cleaned at a tail of a log of data is copied to a head of the log, [Zwilling col. 5, lines 34-52 with Zwilling, Fig. 2 with Zwilling, col. 12, lines 15-17] wherein the range to be cleaned is a range of addresses in the storage system,…stored from each one of a plurality of remote computer systems" [Zwilling, col. 3, lines 15-23]… the operations including:
  - locating blocks of data in the log that are both referenced by at least one other block of data stored within the nodes of the plurality of storage trees and within the range to be cleaned using pruned walking; [Zwilling, col. 12, lines 24-53 with Zwilling, Fig. 2 with Zwilling, col. 8, lines 49-51 with Zwilling, col. 8, lines 60-64 with Zwilling, Figs. 3C and 3E]
  - copying blocks of data associated with child nodes of a current node that are within the range to be cleaned to the head of the log; [Zwilling col. 5, lines 34-52 with Zwilling, Fig. 2]
  - retrieving a block of data associated with the current node, upon determining that a minimum address value among addresses of descendent nodes is

within the range to be cleaned; [Zwilling col. 5, lines 29-52 with Zwilling, Fig. 2]

- designating, as the current node, one of the child nodes of the current node that is an interior node, upon determining that at least one child node is an interior node; [Zwilling, col. 8, lines 46-65] and
- designating, as the current node, an ancestor node of the current node whose descendent nodes are unprocessed; [Zwilling, col. 8, lines 46-65] and
- marking the range as unallocated when the blocks of data that are active and within the range are copied to the head of the log so that at least a portion of the address space within the range to be cleaned can be reclaimed" [Zwilling col. 5, lines 29-52 with Zwilling, Fig. 2].

Zwilling discloses the above limitations but does not expressly teach:

- "…the storage system having a plurality of storage trees representing active snapshots of user-generated data …taken at different points in time, each storage tree having a plurality of nodes, wherein each block of data is represented by a node of the storage tree."

With respect to Claim 15, an analogous art, Hitz, teaches:

- the storage system having a plurality of storage trees representing active snapshots of user-generated data [Hitz, cols. 17-18, lines 66-16] …taken at different points in time, [Hitz, col. 17, lines 40-43] each storage tree having a plurality of nodes, [Hitz, cols. 17-18, lines 66-16 with Zwilling, col. 8, lines 49-64] wherein each block of data is represented by a node of the storage tree"

[Hitz, col. 17, lines 40-49 with Hitz, cols. 17-18, lines 66-16 with Zwilling, col. 8, lines 49-64].

It would have been obvious to one of ordinary skill in the art at the time of invention having the teachings of Hitz and Zwilling before him/her to combine Hitz with Zwilling because both inventions are directed towards conserving file space (Zwilling shrinks and Hitz snapshots data that was changed at a block level (not a fie level)).

Hitz's invention would have been expected to successfully work well with Zwilling's invention because both inventions use file systems on computers. Zwilling discloses an on-line dynamic file shrink facility (title) comprising shrinking log files. However, Zwilling does not expressly disclose storage trees each representing a snapshot and where unreferenced data remains untouched. Hitz discloses a write anywhere file-system layout (title) comprising storage trees (file system snapshot trees) each representing a snapshot and where unreferenced data remains untouched.

It would have been obvious to one of ordinary skill in the art at the time of invention having the teachings of Hitz and Zwilling before him/her to take the write anywhere file-system layout and the storage trees of snapshots from Hitz and install it into the invention of Zwilling, thereby offering the obvious advantage of taking snapshots of garbage collected/shrunken data, thus saving space. Shrinking the files prior to snapshoting them frees more space in Hitz thus creating more room in Hitz so that another later snapshot won't be prematurely forced out when space in Hitz runs out (by snapshots consuming unacceptable numbers of disk blocks). Also, shinking prior to snapshoting guarantees that Hitz will not back up empty blocks (since they are dirtied

by becoming empty in a snapshot). This also conserves space in Hitz. In fact, Hitz desires files to be in a shrunken form (Hitz, col. 5, lines 37-40). Therefore, Hitz is merely using Zwilling's shrinking technique to obtain files that have no fragments.

Even though Zwilling is only one reference being used to reject Claim 15, the rejection on Claim 15 is under 35 U.S.C. 103(a) because different embodiments of Zwilling are use in the rejection for Claim 15 and its respective dependant claims. Zwilling teaches all of Claim 15 through Zwilling's different embodiments.

Claim 16 can be mapped to Zwilling (as modified by Hitz) as follows: "The method of claim 15, wherein performing the following until each block of data that is active in the range to be cleaned at the tail of the log of data is copied to a head of the log includes updating addresses of that copied blocks of data within a location table" [Zwilling, col. 8, lines 21-46 with Zwilling, col. 9, lines 5-7].

Claim 17 can be mapped to Zwilling (as modified by Hitz) as follows: "The method of claim 15, wherein performing the following until each block of data that is active in the range to be cleaned at the tail of the log of data is copied to the head of the log includes updating a minimum address value among addresses of descendent nodes for an entry for the current node in a location table where the minimum address value changes based on copying of the blocks of data associated with the descendent nodes of the current node" [Zwilling, cols. 8-9, lines 21-15 with Zwilling col. 8, lines 20-35 with Zwilling, col. 7, lines 38-21].

Claim 18 can be mapped to Zwilling (as modified by Hitz) as follows: "The method of claim 15, wherein at least one block of data stored in the log is referenced by

more than one of other blocks of data" [Hitz, col. 18, lines 24-30 with Hitz, col. 2, lines 25-36].

For **Claim 27**, Zwilling teaches:

- "…a storage space to store said blocks of data having been allocated in said set of one or more storage devices; [Zwilling, col. 5, lines 29-52 with Zwilling, Fig. 2]
- having stored therein a minimum address value of descendent nodes of interior nodes of said plurality of storage trees; [Zwilling, cols. 8-9, lines 46-15 with Zwilling, col. 9, lines 43-51 with Zwilling col. 8, lines 20-35 with Zwilling, col. 7, lines 38-21] and
- a garbage collection logic to clean a currently selected range from the tail of said log, [Zwilling, col. 5, lines 29-52 with Zwilling, Fig. 2] said garbage collection logic to perform pruned walking of nodes of said plurality of storage trees based on said set of location tables and said currently selected range to locate blocks of data that are referenced by at least one other block of data currently stored within the plurality of storage trees" [Zwilling, cols. 8-9, lines 46-17 with Zwilling, col. 5, lines 29-52 with Zwilling, col. 12, lines 24-53 with Zwilling, Fig. 2 with Zwilling, Figs. 3C and 3E].

With respect to Claim 27 Zwilling teaches in a different embodiment:

- "…stored from a file system located on a remote computer [Zwilling, col. 3, lines 15-23] … stored from the remote computer file system [Zwilling, col. 3, lines 15-23] …from said remote computer file system [Zwilling, col. 3, lines 15-23]…within the remote computer" [Zwilling, col. 3, lines 15-23].

Zwilling discloses the above limitations but does not expressly teach: "A backup system comprising:

- a plurality of storage trees, each storage tree representing an active snapshot of user-generated data … taken at a different time, each storage tree having a plurality of nodes including leaf nodes representing data blocks of the user-generated data … and root and intermediate nodes representing data blocks of system-generated data created by the storage device, the system generated data including pointers to reference other data blocks of the user-generated data to enable overlapping and sharing of the user-generated target data between and among different snapshots of the plurality of snapshots, wherein each leaf node of user-generated data … has been backed up from a set of one or more storage devices
- …from a backup storage space
- …a set of one or more location tables."

With respect to Claim 27, an analogous art, Hitz, teaches: "A backup system [Hitz, col. 17, lines 40-50] comprising:

- a plurality of storage trees, [Hitz, cols. 17-18, lines 66-16] each storage tree representing an active snapshot of user-generated data [Hitz, cols. 17-18, lines 66-16]…taken at a different time, [Hitz, col. 17, lines 40-43] each storage tree having a plurality of nodes [Hitz, cols. 17-18, lines 66-16 with Zwilling, col. 8, lines 49-64] including leaf nodes representing data blocks of the user-generated data [Hitz, cols. 17-18, lines 66-14] … and root and intermediate nodes

representing data blocks of system-generated data created by the storage device, [Hitz, cols. 17-18, lines 66-14] the system generated data including pointers to reference other data blocks of the user-generated data to enable overlapping and sharing of the user-generated target data between and among different snapshots of the plurality of snapshots, [Hitz, cols. 17-18, lines 66-14 with Hitz, col. 18, lines 35-38] wherein each leaf node of user-generated data has been backed up from a set of one or more storage devices [Hitz, cols. 17-18, lines 65-14]

- …from a backup storage space [Hitz, cols. 17-18, lines 65-14]
- …a set of one or more location tables" [Hitz, cols. 17-18, lines 65-14 with Zwilling, col. 10, lines 5-13].

It would have been obvious to one of ordinary skill in the art at the time of invention having the teachings of Hitz and Zwilling before him/her to combine Hitz with Zwilling because both inventions are directed towards conserving file space (Zwilling shrinks and Hitz snapshots data that was changed at a block level (not a fie level)).

Hitz's invention would have been expected to successfully work well with Zwilling's invention because both inventions use file systems on computers. Zwilling discloses an on-line dynamic file shrink facility (title) comprising shrinking log files. However, Zwilling does not expressly disclose storage trees each representing a snapshot and where unreferenced data remains untouched. Hitz discloses a write anywhere file-system layout (title) comprising storage trees (file system snapshot trees) each representing a snapshot and where unreferenced data remains untouched.

It would have been obvious to one of ordinary skill in the art at the time of invention having the teachings of Hitz and Zwilling before him/her to take the write anywhere file-system layout and the storage trees of snapshots from Hitz and install it into the invention of Zwilling, thereby offering the obvious advantage of taking snapshots of garbage collected/shrunken data, thus saving space. Shrinking the files prior to snapshoting them frees more space in Hitz thus creating more room in Hitz so that another later snapshot won't be prematurely forced out when space in Hitz runs out (by snapshots consuming unacceptable numbers of disk blocks). Also, shinking prior to snapshoting guarantees that Hitz will not back up empty blocks (since they are dirtied by becoming empty in a snapshot). This also conserves space in Hitz. In fact, Hitz desires files to be in a shrunken form (Hitz, col. 5, lines 37-40). Therefore, Hitz is merely using Zwilling's shrinking technique to obtain files that have no fragments.

Different embodiments of Zwilling are use in the rejection for Claim 27 and its respective dependant claims. Zwilling (as modified by Hitz) teaches all of Claim 27 with assistance from Zwilling's different embodiments.

Claim 28 can be mapped to Zwilling (as modified by Hitz kath) as follows: "The backup system of claim 27, wherein two different nodes of a same storage tree reference a same node in the same storage tree" [Hitz, col. 18, lines 24-30 with Hitz, col. 2, lines 25-36].

Claim 29 can be mapped to Zwilling (as modified by Hitz) as follows: "The backup system of claim 27, wherein the garbage collection logic is to update references to a node that is within the currently selected range based on an update to an entry in

the set of one or more location tables" [Zwilling, Fig. 3C with Zwilling, col. 8, lines 29-33 with Zwilling, col. 9, lines 4-7].

Claim 30 can be mapped to Zwilling (as modified by Hitz) as follows: "The backup system of claim 27, wherein the garbage collection logic is to prune walking of the nodes of said plurality of storage trees based on the minimum addresses stored in the set of one ore more location tables" [Zwilling, cols. 8-9, lines 46-17 with Zwilling, col. 5, lines 29-52 with Zwilling, col. 12, lines 24-40 with Zwilling, col. 10, lines 5-13].

For **Claim 31**, Zwilling teaches: "An apparatus [Zwilling, cols. 4-5, lines 62-12] comprising:

- …by recording references to blocks of backed up data [Zwilling, cols. 8-9, lines 46-17 with Zwilling, col. 10, lines 5-13]
- …an allocator logic to allocate contiguous blocks of storage space from a log of a backup storage space to store said blocks of backed up data [Zwilling, col. 5, lines 29-52 with Zwilling, col. 12, lines 24-40]
- a garbage collection logic…to clean a currently selected contiguous range from the tail of said log, [Zwilling, col. 5, lines 29-52, with Zwilling, Fig. 2] said garbage collection logic to,
    - walk only those nodes of said plurality of storage trees that possibly identify those of said blocks of data that are stored in said currently selected contiguous range or that possibly are themselves stored in said currently selected contiguous range, [Zwilling, cols. 8-9, lines 46-14 with Zwilling, col. 9, lines 43-50]

- locate blocks of data that are referenced by at least one other block o data stored within the plurality of storage trees, [Zwilling, col. 12, lines 24-53 with Zwilling, Fig. 2 with Zwilling, col. 8, lines 49-51 with Zwilling, col. 8, lines 60-64 with Zwilling, Figs. 3C and 3E] and
- sweep said currently selected contiguous range, [Zwilling, col. 5, lines 29-52, with Zwilling, Fig. 2] copying blocks of data that are referenced and within the range out of the range and marking the range as unallocated so that at least a portion of the address space within the range can be reclaimed" [Zwilling, col. 5, lines 29-52 with Zwilling, Fig. 2].

With respect to Claim 31 Zwilling teaches in a different embodiment:

- "…located on a remote computer [Zwilling, col. 3, lines 15-23]
- stored from said remote computer file system [Zwilling, col. 3, lines 15-23] … stored from the remote computer file system" [Zwilling, col. 3, lines 15-23].

Zwilling discloses the above limitations but does not expressly teach:

- "…a hardware backup system to backup a file system…, said backup file system including:
  - a tracking logic to generate a plurality of storage trees, each storage tree representing active backup snapshots of user generated data stored from said remote computer file system taken at different times stored in a set of one or more storage devices, each storage tree having a plurality of nodes including leaf nodes representing data blocks of the user-generated data stored from the remote computer file system and root and intermediate nodes

representing data blocks of system-generated data created by the storage device, the system-generated data including pointers to reference other data blocks of the user-generated data to enable overlapping and sharing of the user-generated target data between and among different snapshots of the plurality of snapshots;

- …responsive to deletion of one or more of said active backup snapshots."

With respect to Claim 31, an analogous art, Hitz, teaches:

- "…a hardware backup system [Hitz, col. 17, lines 40-50] to backup a file system …, [Hitz, cols. 17-18, lines 65-14] said backup file system including:
  - a tracking logic to generate a plurality of storage trees, each storage tree representing active backup snapshots of user generated data …taken at different times [Hitz, cols. 17-18, lines 65-14] stored in a set of one or more storage devices, [Hitz, cols. 17-18, lines 65-14] each storage tree having a plurality of nodes [Hitz, cols. 17-18, lines 66-16 with Zwilling, col. 8, lines 49-64] including leaf nodes representing data blocks of the user-generated data [Hitz, cols. 17-18, lines 66-14]… and root and intermediate nodes representing data blocks of system-generated data created by the storage device, [Hitz, cols. 17-18, lines 66-14] the system-generated data including pointers to reference other data blocks of the user-generated data to enable overlapping and sharing of the user-generated target data between and among different snapshots of the plurality of snapshots; [Hitz, cols. 17-18, lines 66-14 with Hitz, col. 18, lines 35-38]

- …responsive to deletion of one or more of said active backup snapshots" [Hitz, cols. 17-18, lines 65-14].

It would have been obvious to one of ordinary skill in the art at the time of invention having the teachings of Hitz and Zwilling before him/her to combine Hitz with Zwilling because both inventions are directed towards conserving file space (Zwilling shrinks and Hitz snapshots data that was changed at a block level (not a fie level)).

Hitz's invention would have been expected to successfully work well with Zwilling's invention because both inventions use file systems on computers. Zwilling discloses an on-line dynamic file shrink facility (title) comprising shrinking log files. However, Zwilling does not expressly disclose storage trees each representing a snapshot and where unreferenced data remains untouched. Hitz discloses a write anywhere file-system layout (title) comprising storage trees (file system snapshot trees) each representing a snapshot and where unreferenced data remains untouched.

It would have been obvious to one of ordinary skill in the art at the time of invention having the teachings of Hitz and Zwilling before him/her to take the write anywhere file-system layout and the storage trees of snapshots from Hitz and install it into the invention of Zwilling, thereby offering the obvious advantage of taking snapshots of garbage collected/shrunken data, thus saving space. Shrinking the files prior to snapshoting them frees more space in Hitz thus creating more room in Hitz so that another later snapshot won't be prematurely forced out when space in Hitz runs out (by snapshots consuming unacceptable numbers of disk blocks). Also, shinking prior to snapshoting guarantees that Hitz will not back up empty blocks (since they are dirtied

by becoming empty in a snapshot). This also conserves space in Hitz. In fact, Hitz desires files to be in a shrunken form (Hitz, col. 5, lines 37-40). Therefore, Hitz is merely using Zwilling's shrinking technique to obtain files that have no fragments.

Different embodiments of Zwilling are use in the rejection for Claim 31 and its respective dependant claims. Zwilling (as modified by Hitz) teaches all of Claim 31 with assistance from Zwilling's different embodiments.

Claim 32 can be mapped to Zwilling (as modified by Hitz) as follows: "The apparatus of claim 31, wherein the plurality of storage trees include interior nodes and leaf nodes, [Zwilling, col. 8, lines 10-20 with Zwilling, cols. 8-9, lines 46-14] the interior nodes to include references to other nodes in one or more of the plurality of storage trees, [Zwilling, cols. 8-9, lines 46-14] wherein two different interior nodes of a same tree references a same node in the same tree" [Hitz, col. 18, lines 24-30 with Hitz, col. 2, lines 25-36].

Claims 33 and 35-39 encompass substantially the same scope of the invention as that of Claims 1, 3-7, respectfully, in addition to a machine-readable medium and some instructions for performing the method steps of Claims 1, 3-7, respectfully. Therefore, Claims 33, 35-39 are rejected for the same reasons as stated above with respect to Claims 1, 3-7, respectfully.

For **Claim 40**, Zwilling teaches: "A computer-storage medium that stores instructions, which when executed by a computer, cause said computer to perform operations in a hardware storage device [Zwilling, col. 3, lines 15-23] comprising:

- garbage collecting within a range of addresses in the storage system [Zwilling, col. 12, lines 24-53 with Zwilling, Fig. 2] having"

With respect to Claim 40, Zwilling teaches in a different embodiment:

- "…stored from each one of a plurality of remote computer systems [Zwilling, col. 3, lines 15-23]…stored from each one of the plurality of remote computer systems [Zwilling, col. 3, lines 15-23] …the garbage collecting including:
  - prune walking of the plurality of storage trees to determine active blocks of data within said range, where active blocks of data are those that are referenced by at least one other block of data still in one of the plurality of storage trees, [Zwilling, cols. 8-9, lines 46-17 with Zwilling, col. 5, lines 29-52 with Zwilling, col. 12, lines 24-53 with Zwilling, Fig. 2 with Zwilling, Figs. 3C and 3E] the prune walking including:
    - determining, based on accessing in one of said plurality of storage trees a parent node that has a plurality of descendent nodes, that none of the plurality of descendant nodes are associated with blocks of data within the range; [Zwilling, col. 5, lines 29-52 with Zwilling, col. 12, lines 24-53 with Zwilling, col. 9, lines 56-65] and
    - skipping the walking of the plurality of descendent nodes based on said determining, [Zwilling col. 5, lines 12-16]
- wherein the active blocks determined to be in the range are copied out of the range and the range is marked as unallocated so that at least a portion of the

address space within the range can be reclaimed" [Zwilling col. 5, lines 29-52 with Zwilling, Fig. 2].

Zwilling discloses the above limitations but does not expressly teach:

- "…a plurality of storage trees, each storage tree representing active snapshots of user-generated data … taken at different times and having a plurality of nodes, the plurality of nodes including leaf nodes representing data blocks of the user-generated data … and root an intermediate nodes representing data blocks of system-generated data created by the storage device, the system-generated data including pointers to reference other data blocks of the user-generated data to enable overlapping and sharing of the user-generated target data between and among different snapshots of the plurality of snapshots and having multiple references to the same block of data."

With respect to Claim 40, an analogous art, Hitz, teaches:

- "…a plurality of storage trees, each storage tree representing active snapshots of user-generated data [Hitz, cols. 17-18, lines 66-16] … taken at different times [Hitz, col. 17, lines 40-43] and having a plurality of nodes, [Hitz, cols. 17-18, lines 66-16 with Zwilling, col. 8, lines 49-64] the plurality of nodes including leaf nodes representing data blocks of the user-generated data [Hitz, cols. 17-18, lines 66-14] … and root an intermediate nodes representing data blocks of system-generated data created by the storage device, [Hitz, cols. 17-18, lines 66-14] the system-generated data including pointers to reference other data blocks of the user-generated data to enable overlapping and sharing of the user-generated

target data between and among different snapshots of the plurality of snapshots [Hitz, cols. 17-18, lines 66-14 with Hitz, col. 18, lines 35-38] and having multiple references to the same block of data" [Hitz, cols. 17-18, lines 66-14 with Hitz, col. 18, lines 35-38].

It would have been obvious to one of ordinary skill in the art at the time of invention having the teachings of Hitz and Zwilling before him/her to combine Hitz with Zwilling because both inventions are directed towards conserving file space (Zwilling shrinks and Hitz snapshots data that was changed at a block level (not a fie level)).

Hitz's invention would have been expected to successfully work well with Zwilling's invention because both inventions use file systems on computers. Zwilling discloses an on-line dynamic file shrink facility (title) comprising shrinking log files. However, Zwilling does not expressly disclose storage trees each representing a snapshot and where unreferenced data remains untouched. Hitz discloses a write anywhere file-system layout (title) comprising storage trees (file system snapshot trees) each representing a snapshot and where unreferenced data remains untouched.

It would have been obvious to one of ordinary skill in the art at the time of invention having the teachings of Hitz and Zwilling before him/her to take the write anywhere file-system layout and the storage trees of snapshots from Hitz and install it into the invention of Zwilling, thereby offering the obvious advantage of taking snapshots of garbage collected/shrunken data, thus saving space. Shrinking the files prior to snapshoting them frees more space in Hitz thus creating more room in Hitz so that another later snapshot won't be prematurely forced out when space in Hitz runs out

(by snapshots consuming unacceptable numbers of disk blocks). Also, shinking prior to snapshoting guarantees that Hitz will not back up empty blocks (since they are dirtied by becoming empty in a snapshot). This also conserves space in Hitz. In fact, Hitz desires files to be in a shrunken form (Hitz, col. 5, lines 37-40). Therefore, Hitz is merely using Zwilling's shrinking technique to obtain files that have no fragments.

Different embodiments of Zwilling are use in the rejection for Claim 40 and its respective dependant claims. Zwilling (as modified by Hitz) teaches all of Claim 40 with assistance from Zwilling's different embodiments.

Claims 41-43, 45, and 46's limitation(s) are met by Claims 9-11, 13, and 14's limitation(s), respectfully. Therefore, Claims 41-43, 45, and 46 are rejected for the same reason(s) as stated below with respect to Claims 9-11, 13, and 14, respectfully.

14. Claims 8-11, 13, 14, 20-22, 24-26, and 47-50 are rejected under 35 U.S.C. 103(a) as being unpatentable over U.S. Patent No. 6,249,792 (Zwilling et al.) in view of U.S. Patent No. 5,963,962 (Hitz et al.), further in view of U.S. Patent No. 5,829,005 (Senator).

For **Claim 8**, Zwilling teaches: "A method comprising:

- garbage collecting within a range of addresses in a storage system [Zwilling, col. 12, lines 24-53 with Zwilling, Fig. 2] having"

With respect to Claim 8, Zwilling teaches in a different embodiment:

- "…stored from each one of a plurality of remote computer systems [Zwilling, col. 3, lines 15-23]

- …wherein each block of data is stored within a log … the garbage collecting including: [Zwilling, cols. 8-9, lines 46-17]
    - pruning walking of the plurality of storage trees to determine active blocks of data within said range, where active blocks of data are those that are referenced by at least one other block of data still stored within one of the plurality of storage trees, [Zwilling, cols. 8-9, lines 46-17 with Zwilling, col. 5, lines 29-52 with Zwilling, col. 12, lines 24-53 with Zwilling, Fig. 2 with Zwilling, Figs. 3C and 3E] the pruning walking including:
        - determining, based on accessing in one of said plurality of storage trees a parent node that has a plurality of descendent nodes, that none of the plurality of descendant nodes are associated with blocks of data within the range; [Zwilling, col. 5, lines 29-52 with Zwilling, col. 12, lines 24-53 with Zwilling, col. 9, lines 56-65] and
        - skipping the walking of the plurality of descendent nodes based on said determining, [Zwilling col. 5, lines 12-16]
- wherein the active blocks determined to be in the range are copied out of the range and the range is marked as unallocated so that at least a portion of the address space within the range can be reclaimed" [Zwilling col. 5, lines 29-52 with Zwilling, Fig. 2].

Zwilling discloses the above limitations but does not expressly teach:

- "…a plurality of storage trees representing active snapshots of user-generated data stored from each one of a plurality of remote computer systems taken at

different points in time, each storage tree having a plurality of nodes and having multiple references to a same block of data… as contiguous data in a sequential order and wraps around to the beginning of the log once the end is reached so that the size of the log does not increase beyond its originally created size."

With respect to Claim 8, an analogous art, Hitz, teaches:

- "…a plurality of storage trees representing active snapshots of user-generated data [Hitz, cols. 17-18, lines 66-16] …taken at different points in time, [Hitz, col. 17, lines 40-43] each storage tree having a plurality of nodes [Hitz, cols. 17-18, lines 66-16 with Zwilling, col. 8, lines 49-64] and having multiple references to a same block of data" [Hitz, cols. 17-18, lines 66-14 with Hitz, col. 18, lines 35-38].

With respect to Claim 8, an analogous art, Senator, teaches:

- "…as contiguous data in a sequential order and wraps around to the beginning of the log once the end is reached so that the size of the log does not increase beyond its originally created size" [Senator, col. 2, lines 3-10 with Senator, col. 3, lines 29-39].

It would have been obvious to one of ordinary skill in the art at the time of invention having the teachings of Hitz, Senator, and Zwilling before him/her to combine Hitz and Senator with Zwilling because the inventions are directed towards file operations.

Hitz's and Senator's inventions would have been expected to successfully work well with Zwilling's invention because the inventions use file systems on computers. Zwilling discloses an on-line dynamic file shrink facility (title) comprising shrinking log

files. However, Zwilling does not expressly disclose storage trees each representing a snapshot and where unreferenced data remains untouched, or a circular log file. Hitz discloses a write anywhere file-system layout (title) comprising storage trees (file system snapshot trees) each representing a snapshot and where unreferenced data remains untouched. Senator discloses a circular size-bounded file technique for a computer operating system (title) comprising circular log files.

It would have been obvious to one of ordinary skill in the art at the time of invention having the teachings of Hitz, Senator, and Zwilling before him/her to take the write anywhere file-system layout and the storage trees of snapshots from Hitz and the circular logs from Senator and install it into the invention of Zwilling, thereby offering the obvious advantage of taking snapshots of garbage collected/shrunken data, thus saving space. Shrinking the files prior to snapshoting them frees more space in Hitz thus creating more room in Hitz so that another later snapshot won't be prematurely forced out when space in Hitz runs out (by snapshots consuming unacceptable numbers of disk blocks). Also, shinking prior to snapshoting guarantees that Hitz will not back up empty blocks (since they are dirtied by becoming empty in a snapshot). This also conserves space in Hitz. In fact, Hitz desires files to be in a shrunken form (Hitz, col. 5, lines 37-40). Therefore, Hitz is merely using Zwilling's shrinking technique to obtain files that have no fragments. The circular logs of Senator provide the obvious advantage of having bounded log sizes where the space for the log can be re-used/reallocated upon use (Senator, col. 2, lines 3-10 with Senator, col. 3, lines 29-39).

Even though Zwilling is only one reference being used to reject Claim 8, the rejection on Claim 8 is under 35 U.S.C. 103(a) because different embodiments of Zwilling are use in the rejection for Claim 8 and its respective dependent claims. Zwilling teaches all of Claim 8 through Zwilling's different embodiments.

Claim 9 can be mapped to Zwilling (as modified by Hitz and Senator) as follows: "The method of claim 8, wherein the blocks of data are stored in a log and the range is a segment of the log" [Zwilling, col. 12, lines 15-53 with Zwilling, Fig. 2].

Claim 10 can be mapped to Zwilling (as modified by Hitz and Senator) as follows: "The method of claim 9, wherein the segment is at the tail of the log" [Zwilling col. 5, lines 34-52].

Claim 11 can be mapped to Zwilling (as modified by Hitz and Senator) as follows: "The method of claim 10, wherein the determining is performed by comparing a minimum offset of the plurality of descendent nodes against the range, [Zwilling, cols. 8-9, lines 46-15 with Zwilling, col. 9, lines 43-51 with Zwilling col. 8, lines 20-35] wherein the minimum offset is accessed when walking the parent node and without walking the plurality of descendent nodes" [Zwilling, cols. 8-9, lines 46-15 with Zwilling, col. 9, lines 43-51 with Zwilling col. 8, lines 20-35 with Zwilling, col. 7, lines 38-21].

Claim 13 can be mapped to Zwilling (as modified by Hitz and Senator) as follows: "The method of claim 8, wherein the range is a segment at the tail of a log and said copying is from the said segment at the tail to a segment at the head of the log" [Zwilling col. 5, lines 34-52 with Zwilling, Fig. 2].

Claim 14 can be mapped to Zwilling (as modified by Hitz and Senator) as follows: "The method of claim 8, wherein said copying includes updating addresses of the copied blocks of data within a location table" [Zwilling, Fig. 3C with Zwilling, col. 8, lines 29-33 with Zwilling, col. 9, lines 4-7].

For **Claim 20**, Zwilling teaches: "A system comprising:

- a storage device to store a number of blocks of data, [Zwilling, col. 5, lines 6-12] wherein the blocks of data that are marked as allocated are non-modifiable, [Zwilling, col. 6, lines 5-10] the blocks of data to be stored as a log; [Zwilling, col. 12, lines 15-17] wherein each block of data within the log is stored …and
- a garbage collection logic to locate the blocks of data that are both referenced by at least one other block stored within a set of one or more storage trees, and within a range at a tail of the log using pruned walking, the range representing and address range to be cleaned within an allocated segment of the log, [Zwilling, col. 12, lines 24-53 with Zwilling, Fig. 2 with Zwilling, col. 8, lines 49-51 with Zwilling, col. 8, lines 60-64 with Zwilling, Figs. 3C and 3E]
- wherein the log is implemented in a hierarchical architecture [Zwilling, cols. 8-9, lines 46-26 with Zwilling, col. 12, lines 31-34]
- …and wherein said garbage collection logic is operable to copy the located blocks of data to a head of the log and mark the range to be cleaned as unallocated so that at least a portion of the address space within the range can be reclaimed" [Zwilling, col. 5, lines 29-52 with Zwilling, Fig. 2].

Zwilling discloses the above limitations but does not expressly teach:

- "…having a plurality of storage trees, each representing a snapshot taken at a point in time of target data being processed, each storage tree having a plurality nodes, and each node representing a block of data of an active snapshot associated with each storage tree
- …as contiguous data in a sequential order and wraps around to the beginning of the log once the end is reached so that the size of the log does not increase beyond its originally created size."

With respect to Claim 20, an analogous art, Hitz, teaches:

- "…having a plurality of storage trees, [Hitz, cols. 17-18, lines 66-16] each representing a snapshot taken at a point in time of target data being processed, [Hitz, col. 17, lines 40-49] each storage tree having a plurality nodes, [Hitz, cols. 17-18, lines 66-16 with Zwilling, col. 8, lines 49-64] and each node representing a block of data of an active snapshot associated with each storage tree" [Hitz, col. 17, lines 40-49 with Hitz, cols. 17-18, lines 66-16].

With respect to Claim 20, an analogous art, Senator, teaches:

- "…as contiguous data in a sequential order and wraps around to the beginning of the log once the end is reached so that the size of the log does not increase beyond its originally created size,;" [Senator, col. 2, lines 3-10 with Senator, col. 3, lines 29-39].

It would have been obvious to one of ordinary skill in the art at the time of invention having the teachings of Hitz, Senator, and Zwilling before him/her to combine

Hitz and Senator with Zwilling because the inventions are directed towards file operations.

Hitz's and Senator's inventions would have been expected to successfully work well with Zwilling's invention because the inventions use file systems on computers. Zwilling discloses an on-line dynamic file shrink facility (title) comprising shrinking log files. However, Zwilling does not expressly disclose storage trees each representing a snapshot and where unreferenced data remains untouched, or a circular log file. Hitz discloses a write anywhere file-system layout (title) comprising storage trees (file system snapshot trees) each representing a snapshot and where unreferenced data remains untouched. Senator discloses a circular size-bounded file technique for a computer operating system (title) comprising circular log files.

It would have been obvious to one of ordinary skill in the art at the time of invention having the teachings of Hitz, Senator, and Zwilling before him/her to take the write anywhere file-system layout and the storage trees of snapshots from Hitz and the circular logs from Senator and install it into the invention of Zwilling, thereby offering the obvious advantage of taking snapshots of garbage collected/shrunken data, thus saving space. Shrinking the files prior to snapshoting them frees more space in Hitz thus creating more room in Hitz so that another later snapshot won't be prematurely forced out when space in Hitz runs out (by snapshots consuming unacceptable numbers of disk blocks). Also, shinking prior to snapshoting guarantees that Hitz will not back up empty blocks (since they are dirtied by becoming empty in a snapshot). This also conserves space in Hitz. In fact, Hitz desires files to be in a shrunken form (Hitz, col. 5,

lines 37-40). Therefore, Hitz is merely using Zwilling's shrinking technique to obtain files that have no fragments. The circular logs of Senator provide the obvious advantage of having bounded log sizes where the space for the log can be re-used/reallocated upon use (Senator, col. 2, lines 3-10 with Senator, col. 3, lines 29-39).

Different embodiments of Zwilling are use in the rejection for Claim 20 and its respective dependant claims. Zwilling (as modified by Hitz) teaches all of Claim 20 with assistance from Zwilling's different embodiments.

Claim 21 can be mapped to Zwilling (as modified by Hitz and Senator) as follows: "The system of claim 20, wherein the garbage collection logic is to copy the blocks of data that are referenced to an unallocated address space of the log" [Zwilling, col. 5, lines 29-52 with Zwilling, Fig. 2].

Claim 22 can be mapped to Zwilling (as modified by Hitz and Senator) as follows: "The system of claim 21, wherein the garbage collection logic is to copy the blocks of data that are referenced to a head of the log" [Zwilling, col. 5, lines 29-52 with Zwilling, Fig. 2].

Claim 24 can be mapped to Zwilling (as modified by Hitz and Senator) as follows: "The system of claim 20 wherein at least one of the number of blocks of data are referenced by more than one reference" [Hitz, col. 18, lines 24-30 with Hitz, col. 2, lines 25-36 with Hitz, col. 20, lines 25-35].

Claim 25 can be mapped to Zwilling (as modified by Hitz and Senator) as follows: "The system of claim 20 comprising a location table to include entries associated with interior nodes of a storage tree, [Zwilling, col. 10, lines 5-13] wherein

each entry is to include a minimum value among the addresses of descendent nodes of the associated interior node" [Zwilling, cols. 8-9, lines 46-15 with Zwilling, col. 9, lines 43-51 with Zwilling col. 8, lines 20-35 with Zwilling, col. 7, lines 38-21].

Claim 26 can be mapped to Zwilling (as modified by Hitz and Senator) (as modified by) as follows: "The system of claim 25, wherein the garbage collection logic is to locate the blocks of data that are referenced and within the range at the tail of the log based on the minimum values stored in the entries of the location table" [Zwilling, col. 5, lines 29-40 with Zwilling, col. 8, lines 46-65 with Zwilling, col. 10, lines 5-13].

For **Claim 47**, Zwilling teaches: "A computer-storage medium that stores instructions, which when executed by a computer, cause said computer to perform garbage collection operations in a hardware storage device [Zwilling, col. 5, lines 29-52 with Zwilling, Fig. 2 with Zwilling, col. 3, lines 15-23] comprising."

With respect to Claim 47, Zwilling teaches in a different embodiment:

- "…performing the following operations until each block of data that is active in a range to be cleaned at a tail of a log of data is copied to a head of the log, [Zwilling col. 5, lines 34-52 with Zwilling, Fig. 2 with Zwilling, col. 12, lines 15-17] wherein the range to be cleaned is a range of addresses in a storage system …, and wherein data stored within the log is stored … the operations including:
- locating blocks of data in a log that are both referenced by at least one other block of data stored within the plurality of storage trees, and within the range using pruned walking; [Zwilling, col. 12, lines 24-53 with Zwilling, Fig. 2 with

Zwilling, col. 8, lines 49-51 with Zwilling, col. 8, lines 60-64 with Zwilling, Figs. 3C and 3E]

- copying blocks of data associated with child nodes of a current node that are within the range to be cleaned to the head of the log; [Zwilling col. 5, lines 34-52 with Zwilling, Fig. 2]
- retrieving a block of data associated with the current node, upon determining that a minimum address value among addresses of descendent nodes is within the range to be cleaned; [Zwilling col. 5, lines 29-52 with Zwilling, Fig. 2]
- designating, as the current node, one of the child nodes of the current node that is an interior node, upon determining that at least one child node is an interior node; [Zwilling, col. 8, lines 46-65]
- designating, as the current node, an ancestor node of the current node whose descendent nodes are unprocessed; [Zwilling, col. 8, lines 46-65] and
- marking the range at the tail of the log as unallocated when the blocks of data that are active and within the range are copied to the head of the log so that at least a portion of an address space within the range can be reclaimed" [Zwilling col. 5, lines 29-52 with Zwilling, Fig. 2].

Zwilling discloses the above limitations but does not expressly teach:

- "…having a plurality of storage trees, each storage tree having a plurality of nodes, wherein a block of data is associated with a nodes of the storage tree

- …as contiguous data in a sequential order and wraps around to the beginning of the log once the end is reached so that the size of the log does not increase beyond its originally created size."

With respect to Claim 47, an analogous art, Hitz, teaches:

- "…having a plurality of storage trees, [Hitz, cols. 17-18, lines 66-16] each storage tree having a plurality of nodes, [Hitz, cols. 17-18, lines 66-16] wherein a block of data is associated with a nodes of the storage tree" [Hitz, col. 17, lines 40-49 with Hitz, cols. 17-18, lines 66-16 with Zwilling, col. 8, lines 49-64].

With respect to Claim 47, an analogous art, Senator, teaches:

- "…as contiguous data in a sequential order and wraps around to the beginning of the log once the end is reached so that the size of the log does not increase beyond its originally created size" [Senator, col. 2, lines 3-10 with Senator, col. 3, lines 29-39].

It would have been obvious to one of ordinary skill in the art at the time of invention having the teachings of Hitz, Senator, and Zwilling before him/her to combine Hitz and Senator with Zwilling because the inventions are directed towards file operations.

Hitz's and Senator's inventions would have been expected to successfully work well with Zwilling's invention because the inventions use file systems on computers. Zwilling discloses an on-line dynamic file shrink facility (title) comprising shrinking log files. However, Zwilling does not expressly disclose storage trees each representing a snapshot and where unreferenced data remains untouched, or a circular log file. Hitz

discloses a write anywhere file-system layout (title) comprising storage trees (file system snapshot trees) each representing a snapshot and where unreferenced data remains untouched. Senator discloses a circular size-bounded file technique for a computer operating system (title) comprising circular log files.

It would have been obvious to one of ordinary skill in the art at the time of invention having the teachings of Hitz, Senator, and Zwilling before him/her to take the write anywhere file-system layout and the storage trees of snapshots from Hitz and the circular logs from Senator and install it into the invention of Zwilling, thereby offering the obvious advantage of taking snapshots of garbage collected/shrunken data, thus saving space. Shrinking the files prior to snapshoting them frees more space in Hitz thus creating more room in Hitz so that another later snapshot won't be prematurely forced out when space in Hitz runs out (by snapshots consuming unacceptable numbers of disk blocks). Also, shinking prior to snapshoting guarantees that Hitz will not back up empty blocks (since they are dirtied by becoming empty in a snapshot). This also conserves space in Hitz. In fact, Hitz desires files to be in a shrunken form (Hitz, col. 5, lines 37-40). Therefore, Hitz is merely using Zwilling's shrinking technique to obtain files that have no fragments. The circular logs of Senator provide the obvious advantage of having bounded log sizes where the space for the log can be re-used/reallocated upon use (Senator, col. 2, lines 3-10 with Senator, col. 3, lines 29-39).

Different embodiments of Zwilling are use in the rejection for Claim 20 and its respective dependant claims. Zwilling (as modified by Hitz) teaches all of Claim 20 with assistance from Zwilling's different embodiments.

Claims 48-50's limitation(s) have already been met by Claims 16-18's limitation(s), respectfully. Therefore, Claims 48-50 are rejected for the same reason(s) as stated above with respect to Claims 16-18, respectfully.

Conclusion

15. Any prior art made of record and not relied upon is considered pertinent to applicant's disclosure. Applicant is advised that, although not used in the rejections above, prior art cited on any PTO-892 form and not relied upon is considered materially relevant to the applicant's claimed invention and/or portions of the claimed invention.

Any inquiry concerning this communication or earlier communications from the examiner should be directed to Brent S. Stace whose telephone number is 571-272-8372 and fax number is 571-273-8372. The examiner can normally be reached on M-F 9am-5:30pm EST.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Apu M. Mofiz can be reached on 571-272-4080. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see http://pair-direct.uspto.gov. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free).

/B. S./

Examiner, Art Unit 2161

/Apu M Mofiz/
Supervisory Patent Examiner, Art Unit 2161